



PATENT
CISCO-8384

This Application is submitted in the names of Inventors Hans Henrik Bostrom, Senthil Arunachalam, Sze-Wa Lao, Srinath Venkatesan to Cisco Technologies, Inc. a California Corporation.

SPECIFICATION

METHOD FOR PROVIDING PROTOCOL AGGREGATION AS AN END-TO-END SERVICE ACROSS A TUNNELING NETWORK

FIELD OF THE INVENTION

[0001] The present invention relates broadly to a computer network serving data transmission functions between geographically separated customer sites. More specifically, the present invention relates to providing data transmission connections end-to-end across an internet service provider computer network.

BACKGROUND OF THE INVENTION

[0002] Business customers of Internet service providers (ISPs) desire transparent tunneling of port aggregation protocol (PAgP) packets through the use of Layer 2 protocol tunneling (L2PT) techniques. This need arises from the need to connect ports of a same Etherchannel between different edge switches of the ISP network to increase availability. Such functionality would allow emulation of

point-to-point connection of Etherchannel ports through an ISP network without dedicated lines, thus maximizing network usage.

[0003] However, current applications lack flexibility. For example, tunneling of link aggregation control protocol (LACP), and unidirectional link detection (UDLD) is not supported. In typical ISP networks, the point-to-point nature of the protocol semantics of UDLD, PAgP and LACP present L2PT requirements that are not supported by prior-art L2PT frameworks that are designed for handling multicast types of protocols, such as spanning tree protocol (STP).

[0004] L2PT is built on top of ISP access VLAN tagging schemes. ISP access VLAN tagging enables ISPs to segregate traffic to and from different customers in the ISP's infrastructure while the customers may appear to be on the same VLANs. Interfaces on the customer networks directed toward the ISP are typically configured as 802.1Q trunks and the interfaces on the edge switches towards the customer networks are typically configured as non-trunking interfaces to create an asymmetric link. Asymmetric links on the ISP side are configured with access VLANs that are unique to each customer.

[0005] Typically, the frames coming out from the customer networks are 802.1Q tagged with appropriate VLAN information. Each tag is preserved when

the frame enters the edge switches in the ISP cloud. Upon exiting the edge switches into the ISP cloud, the frame is double-tagged with outer tag identifying the customer's access VLAN and the inner VLAN identifying the VLAN of the incoming traffic. The core switches inside the ISP infrastructure forward the traffic as normal Ethernet-tagged frames. Upon exiting the edge switch from the ISP towards the customer networks, the outer tag is stripped and the frame is sent out as regular 802.1Q tagged frame so that the original VLAN numbers in the customer networks are recovered. If traffic coming out of the customer network is not tagged, such as in the case of native VLAN frames, then these packets are bridged or routed as if they were normal packets. These packets entering the ISP infrastructure have only a single tag, which is a customer-specific access VLAN tag.

[0006] The above concept can be extended in a hierarchical manner, with border switches performing the double-tagging and the innermost switches in the ISP infrastructure performing the multiple tagging. The amount of multiple tagging that can be supported depends on the maximum length of the Ethernet frame that the hardware can handle.

[0007] L2PT allows switches on the inbound side of the ISP infrastructure to encapsulate protocol packets with a special MAC address and send them across the ISP infrastructure. Edge switches on the outbound side of the ISP

infrastructure decapsulate the protocol packets and send them to a customer network. Core switches inside the ISP infrastructure forward these packets as normal packets without performing any processing on the packets. Thus, the ISP infrastructure is transparent to the customer network.

[0008] When Layer 2 protocol packets enter the tunnel port on the inbound switch, the destination MAC address of the L2 protocol packet is replaced with tunnel MAC (0100.0ccd.cdd0) and forwarded to all trunk ports. These packets also are double-tagged with an outer tag being the customer's access VLAN tag and the inner tag being the customer-specific VLAN tag. The core switches forward these packets to all the ports in the same metro VLAN without performing any processing on the packets. The edge switches on the outbound side restore the proper L2 protocol MAC and forward them to tunnel ports in the same metro VLAN. Thus, the L2 protocol packets are kept intact and delivered to other side of the customer network across the ISP infrastructure.

[0009] Referring to FIG. 1, Corp A and Corp X are connected to access VLANs R and G, respectively, and asymmetric links are created towards the networks in Site1. The Layer 2 protocol packets (for example LACP packets) coming into S1 from Corp A in Site 1 are forwarded into the ISP infrastructure as double-tagged packets with a destination MAC address set to the tunnel MAC address. These double-tagged packets have the outer VLAN tag as R and the

inner tag as (for example) 100. When these double-tagged frames reaches S3, the outer VLAN tag R is removed, the tunnel MAC address is replaced with the respective Layer 2 protocol MAC and the frames are sent out to Corp A on Site 2 as single-tagged frames (i.e. with a tag reading VLAN 100).

[0010] The ISP access VLAN tagging scheme also can be enabled in access ports. In such an application, the encapsulation and de-encapsulation behavior are the same as above, except that these packets in the ISP aren't double-tagged. The single tag is customer-specific access VLAN tag.

SUMMARY

[0011] In one aspect, the present invention provides a method of providing data transmission across a computer network. The method comprises creating a plurality of tunnels across a computer network to connect a first computer to a second computer, where the plurality of tunnels includes a tunnel for each link in a link aggregation of a plurality of transmission protocols. A connection is established through the computer network between a first computer at a first site with a second computer at a second site using the plurality of tunnels. Packets are transmitted end-to-end from the first computer to the second computer in a manner characterized that the computer network preserves a connection from the first computer to the second computer without terminating the connection at the inbound edge of the computer network, with the packets conforming to protocols

in the plurality of transmission protocols. In embodiments, the plurality of transmission protocols includes the LACP and PAgP protocols, and packets are transmitted in accordance with these protocols to perform Ethernet loadsharing across multiple links. In an embodiment, the plurality of transmission protocols comprises the UDLD protocol, and packets are transmitted in accordance with the UDLD protocol to perform unidirectional link detection. A unique ISP access VLAN is assigned to each connection between corresponding Etherchannel ports. In an embodiment, the computer network of the present invention is configured to detect multipoint protocol tunneling. The monitoring can be performed on a per-interface basis, a per-protocol basis, or a per-port basis. The monitoring can be performed by examining a source media access control address on a transmitted protocol data unit. The source media access control address is recorded as a multipoint protocol tunneling reference, and an aging timer is set to a minimum time that is longer than a longest expected transmission time for the transmitted protocol data unit. Before expiration of the aging timer, all packets arriving with a source media access control address other than the reference are dropped. After expiration of the aging timer, the first packet arriving after expiration of the aging timer provides its source media access control address as the next multipoint protocol tunneling reference.

[0012] Embodiments of the present invention include computer program products and network systems that encompass the functionality of the present

invention. Many other features and advantages of the present invention will be realized from reading the following detailed description in conjunction with the drawings, in which:

BRIEF DESCRIPTION OF THE DRAWINGS

[0013] FIG. 1 illustrates a prior-art ISP network configuration; and

[0014] FIG. 2 illustrates a network in accordance with the present invention.

DETAILED DESCRIPTION

[0015] Directing attention to FIG. 2, in an embodiment, the present invention creates a plurality of separate tunnels 10 through an ISP's network 20, one for each link in the UDLD, IEEE 802.ad and/or PAgP link aggregation, and transport LACP, PAgP and/or UDLD packets to opposite ends of the tunnel using L2 tunneling protocol techniques. In an embodiment, LACP, PAgP protocol packets are used for Ethernet load sharing across multiple links, whereas UDLD is used for unidirectional link detection. These protocols play a vital role from an ISP customer 30's point of view when they are deployed. Using the present invention, ISPs can provide its customers with end to end UDLD, IEEE 802.3as and/or PAgP link aggregation using existing ISP infrastructure previously used for providing conventional transparent LAN service where link aggregation is used to access the transparent LAN service.

[0016] Performing data transmission on a tunneled protocol basis allows more flexibility in terms of maintaining point-to-point connections. In embodiments of the present invention, both LACP and UDLD are enabled as tunneled protocols. In an embodiment, a timer is maintained on per-protocol basis. If for some reason if transmission using the UDLD protocol crashes, the LACP connection can be maintained, unlike in a per-port basis.

[0017] Creating virtual cables through a network provides a transparent LAN service. A significant challenge overcome by the present invention is to process the bundle of connections from one side of an ISP network to the other. Normally, in prior-art applications, the Etherchannel connection is terminated at one end of an ISP's network. For example, if a business customer is connected to an ISP, and wants to connect via the ISP to same business located at another site, the client prefers to form a tunnel across the ISP to the other site with the appearance of being a private network. In typical ISPs, the connection the customer is trying to make on the first hop (the ISP side) is terminated on the ISP side and ISP switches determine how the data is routed.

[0018] In contrast, the present invention preserves the connections in separate channels from end to end across the ISP network. The present invention provides a point-to-point topology within the ISP network, emulating cables that connect individual Etherchannel member ports on customer switches across the

ISP network. UDLD, PAgP and LACP exchange protocol data units (PDUs) between each pair of customer switch ports in the Etherchannel through these virtual cables. Once an Etherchannel has been formed between two customer switches across the ISP network, the dynamic load distribution function of the Etherchannel protects the availability of the transparent LAN service between the two customer switches against equipment failures within the ISP network which might result in breaking one or more of the virtual cables connecting the Etherchannel ports. Note that the present invention is not to be confused with the formation of Etherchannels between a customer switch and the ISP edge switch in which there is no L2PT involved.

[0019] The present invention provides a method of network provisioning to enable a point-to-point connection across an ISP network between two corresponding ports in an Etherchannel. This aspect of the present invention also results in network operation that safeguards against error conditions within the ISP network that violates the point-to-point topology required by UDLD, PAgP and LACP. Another benefit is that the present invention results in network operation that dynamically adapts to changes in a customer network's Etherchannel configuration.

[0020] The present invention assigns a unique ISP access VLAN to each virtual cable that connects the corresponding Etherchannel ports. For example, for

a four member Etherchannel, the customer must be assigned four unique ISP access VLANs in order to segregate the point-to-point protocol traffic within each “virtual” cable. This method of network provisioning differs significantly from that of typical L2PT network provisioning schemes for multicast protocols where each customer is assigned only one unique ISP access VLAN.

[0021] Erroneous network configuration that allows tunneled point-to-point protocol packets to be delivered to many points can lead to serious network issues in the ISP customer’s network. Point-to-point tunneled protocols are designed to work between two end points. The presence and participation of a third end point in the protocol interaction can result in the entering an error state in the tunneled protocol’s finite state machine, which can in turn lead to disabling of the tunneled protocol together. In the case of LACP or PAgP, this could lead to inability to perform bundling of Etherchannel ports. To decrease link down detection time, UDLD is enabled whenever tunneling of PAgP or LACP is configured.

[0022] In an embodiment, the present invention provides a mechanism to detect problems due to multipoint protocol tunneling (MPT). MPT detection is useful to prevent the serious issues described above that can arise when PDUs are sent to multiple end points. In an embodiment, MPT detection is enabled/disabled on a per-interface basis. If the user configures a point-to-point tunneled L2 protocol, MPT detection is automatically enabled. A user can also manually

configure MPT detection. When MPT detection is enabled, the system monitors the source media access control (MAC) address of the packets coming out the L2PT tunnel. The source MAC address of the first tunneled PDU from a peer is recorded as an MPT detection reference, and an aging timer for the MPT detection reference is started. The aging timer is maintained in various embodiments on a per-port basis or a per-tunneled protocol basis. In an embodiment, the aging timer is set to a minimum time that is longer than the longest expected time for transmission of a PDU. A multiplier (such as twice the longest expected time for transmission of a PDU) can be used to allow for at least one retransmission of a PDU. In an embodiment, the multiplier is configurable by a user. If packet loss is unlikely to occur, then any value longer than longest expected tunneled protocol interval period is sufficient to safeguard against transient network conditions under which the tunneled PDUs might be dropped. Once a remote end point has been accepted as the tunneled protocol session partner for the local end point, that remote end point is considered to be the legitimate tunneled protocol session partner unless the remote end point does not comply with the tunneled protocol's transmit requirement consistently, for example, for at least twice the period of the longest tunneled protocol interval. While the MPT detection reference is not aged out (the local end point does not receive any tunneled protocol packets from the established remote end point during an aging timer period), all packets arriving with a source MAC address that differ from the MPT detection reference are dropped. In case the MPT detection reference is aged out, the source MAC

address of the next valid packet is used as the new MPT detection reference. Thus the present invention also enables the customer switch to dynamically adapt to Etherchannel configuration changes in the customer network. A user interface is provided to allow the user to manually delete the MPT detection reference. For every n number of MPT detections, where n is a user-configurable MPT detection threshold, an error message is presented to the user and/or recorded in a system log.

[0023] Performing data transmission on a tunneled protocol basis allows more flexibility in terms of maintaining point-to-point connections. In embodiments of the present invention, both LACP and UDLD are enabled separately and in combination as tunneled protocols. In an embodiment, the MPT timer is maintained on per-protocol basis. If for some reason if transmission using the UDLD protocol crashes, the LACP connection can be maintained, unlike in a per-port basis.

[0024] Using a protocol-based MPT timer in accordance with the present invention, for example one timer for each active tunneled protocol on the port, provides more visibility into a system run-time operation. While this embodiment of the present invention provides the ability to track time out event even for tunneled protocols with shorter protocol intervals, but requires more system resources to implement a protocol-based MPT timer.

[0025] A port-based MPT timer may also be used in accordance with the present invention for minimizing resource requirements. However, since the longest protocol interval is used for the port-based timer, if the L2PT transmission is disabled for that protocol with a port-based timer the time out period needs to be adjusted to the next longest protocol interval in order to provide good system response throughout the network.

[0026] While embodiments of the present invention have been illustrated and described in detail, it is to be understood that many modifications can be made to various embodiments of the present invention without departing from the spirit thereof.